Interference Searched EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	14	((miss or mishit) near2 (log\$1 or history or table\$1 or list\$1)).clm.	US-PGPUB	OR	OFF	2007/02/02 13:40
L4	7	(((miss or mishit) near2 (log\$1 or history or table\$1 or list\$1)) same cache).clm.	US-PGPUB	OR	OFF	2007/02/02 13:42
L5	7	(((refresh\$3 or updat\$3) near5 (data adj object)) with cache).clm.	US-PGPUB	OR	OFF	2007/02/02 13:43
L6	8	(((refresh\$3 or updat\$3) same (discard\$3 or LRU) same (predict\$3 or pre-dict\$3 or preload\$3 or pre-load\$3 or prefetch\$3 or pre-fetch\$3))).clm.	US-PGPUB	OR	OFF	2007/02/02 13:43
L7	2	(((discard\$3 or delet\$3 or remov\$3) near2 (list or queue or FIFO)) same ((refresh\$3 or updat\$3) near2 (list or queue or FIFO)) same cache).clm.	US-PGPUB	OR	OFF	2007/02/02 13:44

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	14	((miss or mishit) near2 (log\$1 or history or table\$1 or list\$1)).clm.	US-PGPUB	OR	OFF	2007/02/02 13:40
L4	7	(((miss or mishit) near2 (log\$1 or history or table\$1 or list\$1)) same cache).clm.	US-PGPUB	OR	OFF .	2007/02/02 13:42
L5	7	(((refresh\$3 or updat\$3) near5 (data adj object)) with cache).clm.	US-PGPUB	OR	OFF	2007/02/02 13:43
L6	8	(((refresh\$3 or updat\$3) same (discard\$3 or LRU) same (predict\$3 or pre-dict\$3 or preload\$3 or pre-load\$3 or prefetch\$3 or pre-fetch\$3))).clm.	US-PGPUB	OR	OFF	2007/02/02 13:43
L7	2	(((discard\$3 or delet\$3 or remov\$3) near2 (list or queue or FIFO)) same ((refresh\$3 or updat\$3) near2 (list or queue or FIFO)) same cache).clm.	US-PGPUB	OR	OFF	2007/02/02 13:44
L8	572	(miss or mishit) near2 (log\$1 or history or table\$1 or list\$1)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR .	OFF	2007/02/02 13:45
L9	170	((miss or mishit) near2 (log\$1 or history or table\$1 or list\$1)) same cache	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 13:46
L10	38	((miss or mishit) near2 (log\$1 or history or table\$1 or list\$1)) same cache same (prefetch\$3 or predict\$3 or preload\$3)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR,	OFF	2007/02/02 13:46
L11	1	((miss or mishit) near2 (log\$1 or history or table\$1 or list\$1)) same cache same (prefetch\$3 or predict\$3 or preload\$3) same LRU	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 13:46

L12	12	((miss or mishit) near2 (log\$1 or history or table\$1 or list\$1)) same cache same (prefetch\$3 or predict\$3 or preload\$3) and LRU	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 13:47
L13	108	((miss or mishit) near2 (log\$1 or history or table\$1 or list\$1)) near5 cache	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 13:47
L14	62	(cache near10 (miss or mishit) near2 (log\$1 or history or table\$1 or list\$1)) same (prefetch\$3 or predict\$3 or preload\$3 or pre-fetch\$3 or pre-load\$3)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 13:47
L15	16	((miss or mishit) near2 (log\$1 or history or table\$1 or list\$1)) same LRU	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 13:47
L16	6	(cache near10 (miss or mishit) near2 (log\$1 or history or table\$1 or list\$1)) same (prefetch\$3 or predict\$3 or preload\$3 or pre-fetch\$3 or pre-load\$3) same (discard\$3 or LRU)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 13:48
S1	557	(miss or misthit) near2 (log\$1 or history or table\$1 or list\$1)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 13:45
S2	162	((miss or misthit) near2 (log\$1 or history or table\$1 or list\$1)) same cache	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 13:46

S3	37	((miss or misthit) near2 (log\$1 or history or table\$1 or list\$1)) same cache same (prefetch\$3 or predict\$3 or preload\$3)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 13:46
54	1	((miss or misthit) near2 (log\$1 or history or table\$1 or list\$1)) same cache same (prefetch\$3 or predict\$3 or preload\$3) same LRU	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 13:46
S5	12	((miss or misthit) near2 (log\$1 or history or table\$1 or list\$1)) same cache same (prefetch\$3 or predict\$3 or preload\$3) and LRU	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 13:47
S6	100	((miss or misthit) near2 (log\$1 or history or table\$1 or list\$1)) near5 cache	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 13:47
S7		((miss or misthit) near2 (log\$1 or history or table\$1 or list\$1)) same LRU	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 13:47
S8	61	(cache near10 (miss or misthit) near2 (log\$1 or history or table\$1 or list\$1)) same (prefetch\$3 or predict\$3 or preload\$3 or pre-fetch\$3 or pre-load\$3)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 13:47
S9		(cache near10 (miss or misthit) near2 (log\$1 or history or table\$1 or list\$1)) same (prefetch\$3 or predict\$3 or preload\$3 or pre-fetch\$3 or pre-load\$3) same (discard\$3 or LRU)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 13:48

	<del>,</del>	T		1	·	
S10	71	((refresh\$3 or updat\$3) near5 (data adj object)) with cache	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 13:42
S11	1	((refresh\$3 or updat\$3) near5 (data adj object) near10 (policy or rule\$1)) with cache	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 10:34
S12	17579	((refresh\$3 or updat\$3) and (discard\$3 or LRU) and (predict\$3 or pre-dict\$3 or preload\$3 or pre-load\$3 or prefetch\$3 or pre-fetch\$3))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR .	OFF	2007/02/02 10:36
S13	557	((refresh\$3 or updat\$3) same (discard\$3 or LRU) same (predict\$3 or pre-dict\$3 or preload\$3 or pre-load\$3 or prefetch\$3 or pre-fetch\$3))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 13:43
S14	357	((refresh\$3 or updat\$3) same (discard\$3 or LRU) same (predict\$3 or pre-dict\$3 or preload\$3 or pre-load\$3 or prefetch\$3 or pre-fetch\$3)) same cache	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 10:36
S15	44	((refresh\$3 or updat\$3) same (discard\$3 or LRU) same (predict\$3 or pre-dict\$3 or preload\$3 or pre-load\$3 or prefetch\$3 or pre-fetch\$3)) same (cache near10 (miss or misshit))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR .	OFF	2007/02/02 12:48
S16	1277	((discard\$3 or delet\$3 or remov\$3) near2 (list or queue or FIFO)) same ((refresh\$3 or updat\$3) near2 (list or queue or FIFO))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 11:44

				ı—-		
S17	34	((discard\$3 or delet\$3 or remov\$3) near2 (list or queue or FIFO)) same ((refresh\$3 or updat\$3) near2 (list or queue or FIFO)) same cache	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 13:44
S18	0	((discard\$3 or delet\$3 or remov\$3) near2 (list or queue or FIFO)) same ((refresh\$3 or updat\$3) near2 (list or queue or FIFO) near10 (concurrent\$2 or simultanous\$2)) same cache	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 10:51
S19	3	((discard\$3 or delet\$3 or remov\$3) near2 (list or queue or FIFO)) same ((refresh\$3 or updat\$3) near2 (list or queue or FIFO) near10 (concurrent\$2 or simultanous\$2))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 10:55
S20	704	(711/137).CCLS.	USPAT; USOCR	OR	OFF	2007/02/02 10:55
S21	751	(711/133).CCLS.	USPAT; USOCR	OR ·	OFF	2007/02/02 10:56
S22	316	(711/135).CCLS.	USPAT; USOCR	OR	OFF .	2007/02/02 10:56
S23	375	(711/136).CCLS.	USPAT; USOCR	OR	OFF	2007/02/02 10:56
S24 ·	532	(711/168).CCLS.	USPAT; USOCR	OR	OFF ·	2007/02/02 10:56
S25	9	S6 and S20	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 11:50
S26	6	S6 and S21	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 11:44

					•	
S27	2	S6 and S22	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 11:44
S28 _	5	S6 and S23	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 11:44
S29	-	S6 and S24	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 11:45
S30	1493	(711/118).CCLS.	USPAT; USOCR	OR	OFF	2007/02/02 11:45
S31	226	(711/204).CCLS.	USPAT; USOCR	OR	OFF	2007/02/02 11:45
S32	445	(711/213).CCLS.	USPAT; USOCR	OR	OFF	2007/02/02 11:45
S33	440	(712/207).CCLS.	USPAT; USOCR	OR	OFF	2007/02/02 11:45
S34	18	S6 and S30	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 11:50
S35	0	S6 and S31	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 11:45
S36	5 ·	S6 and S32	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 11:46

S37	2	S6 and S33	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR .	OFF	2007/02/02 11:46
S38	· 34	S14 and S20	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 11:50
S39	11	S14 and S21	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 11:50
S40		S14 and S22	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 11:50
S41	9	S14 and S23	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 11:50
S42	0	S14 and S24	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 11:50
S43	26	S14 and S30	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 11:50

S44	14	S14 and S31	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 11:50
S45	16	S14 and S32	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 11:50
S46	17	S14 and S33	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 12:39
S47	411	(cache near10 object\$1 near10 (miss\$2 or mishit))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 12:40
S48	27	(cache near5 (log or list or table or history) near10 object\$1 near10 (miss\$2 or mishit))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 12:41
S49	44	((refresh\$3 or updat\$3) same (discard\$3 or LRU) same (predict\$3 or pre-dict\$3 or preload\$3 or pre-load\$3 or prefetch\$3 or pre-fetch\$3 or lookahead or look-ahead)) same (cache near10 (miss or misshit))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 12:48
S50	0	((refresh\$3 or updat\$3) same (discard\$3 or LRU) same (predict\$3 or pre-dict\$3 or preload\$3 or pre-load\$3 or prefetch\$3 or pre-fetch\$3 or lookahead or look-ahead)) same (cache near10 (miss or misshit)) same object\$1	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/02/02 12:49



Home | Login | Logout | Access Information | Alerts |

#### Welcome United States Patent and Trademark Office

Search Results

BROWSE

SEARCH

IEEE XPLORE GUIDE

Results for "(cache log, data object, discard queue<and>predict)<and>refresh queue"

☑ e-mail

Your search matched 0 documents. A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

» Search Options

View Session History

**Modify Search** 

New Search

(cache log, data object, discard queue<and>predict)<and>refresh queue

Search

» Кеу

Check to search only within this results set

ieee jnl

IEEE Journal or

Magazine

IEE JNL

IEE Journal or Magazine

IEEE CNF

**IEEE Conference** 

Proceeding

IEE ONF

IEE Conference

Proceeding

IEEE STO IEEE Standard

No results were found.

Please edit your search criteria and try again. Refer to the Help pages if you need assistan

search.

Contact Us Privacy &:

© Copyright 2008 IEEE ---

indexed by 🗓 inspec



Home | Login | Logout | Access Information

### Welcome United States Patent and Trademark Office

BROWSE

SEARCH

IEEE XPLORE GUIDE

0	OPTION 1 Enter keywords or phrases, select field	is, and select operators	(2) Help	» Publications • Select publications
	cache log, discard queue	In Full Text & All Fields	•	IEEE Periodicals
	AND refresh queue	in Full Text & All Fields	•	▼ IEE Periodicals ▼ IEEE Conference
	AND •	In Full Text & All Fletds	•	IEE Conference P IEEE Standards
				₹ IEEE Standards
	» Note: If you use all three search boxes, take precedence over the entry in the third		kes	» Other Resources (Availat
0	OPTION 2 Enter keywords, phrases, or a Boolean	expression	② Help	» Select date range  © Search latest content u  © From year All  to Present
				» Display Format
				© Citation Citatio
	» Note: You may use the search operator	s <and> or <or></or></and>		» Organize results
	without the start and end brackets <>.  » Learn more about <u>Field Codes</u> , <u>Search</u>	Examples, and Search Opera	ators	Maximum 100
				Display 25 resu
				Sort by Relevance
				In Descending
				· Help Contact Us

inspec

© Copyright 29



Home | Login | Logout | Access Information | Aleris |

#### Welcome United States Patent and Trademark Office

Search Re	sults		BROW	/SE	SEARCH	IEEE	XPLORE G	JIDE
Your searc	"(cache log, discard que h matched 28 of 1484991 n of 100 results are display	documents.		elevance in	<b>Descending</b> ord	ler.		⊠e-mail
» Search O	pilons	Modify	/ Search				. •	٠
View Sessi	on History	(cache	log, discard queue	<and>refresh</and>	queue)			Search
New Searc	h	Ch	eck to search onl	y within this	results set			
» Key		Displa	y Format: 🌘	Citation (	Citation & Ab	stract		
IEEE JNL	IEEE Journal or Magazine	← view	selected items	Select A	II Deselect All			
IEE JNL	IEE Journal or Magazine							
IEEE CNF	IEEE Conference Proceeding		Shudong Jin; E	Bestavros, A	; Iyengar, A.;	:		
IEE CNF	IEE Conference Proceeding		2-5 July 2002 F	Page(s):153	tems, 2002, Pro - 160 1109/ICDCS.200	_		ioriai Comei
IEEE STO	IEEE Standard			Full Text: Pt	<u>DF(</u> 356 KB) IE		_	
			Digital Object I	sue 3, May- dentifier 10. Full Text: <u>P</u> [	/ <b>on cache resp</b> June 2000 Page 1109/65.844497 <u>DF(</u> 580 KB) IE	e(s):24 - 28		
	·, :	, m	Selected Areas Volume 18, Iss Digital Object I	att, I.; Clark, s in Commur sue 12, Dec dentifier 10. References	<b>bl</b> C.; Crosby, S.; nications. IEEE J c. 2000 Page(s): 1109/49.898743   Full Text: <u>PDF</u> (	2629 - 263		
		<b>4</b>	International Volume 2, 26-2 Digital Object lo	Park, E.K.; H ware and Ap 28 July 2005 dentifier 10.	an, Y.; oplications Confe 5 Page(s):37 - 38 1109/COMPSAC	rence, 200 3 Vol. 1 3.2005.46		
	*		AbstractPlus   I Rights and Per		<u>)F(</u> 68 KB) 1EE	e Cnf		
		<u> </u>	. Adaptive load	balancing of	content address	s hashing	routing for r	everse pro

Takenaka, T.; Kato, S.; Okamoto, H.;

Communications, 2004 IEEE International Conference on Volume 3, 20-24 June 2004 Page(s):1522 - 1526 Vol.3 Digital Object Identifier 10.1109/ICC.2004.1312765

AbstractPlus | Full Text: PDF(394 KB) IEEE CNF Rights and Permissions 6. A component-based infrastructure for customized persistent object mana Garcia-Banuelos, L.; Duong, P.-Q.; Collet, C.; Database and Expert Systems Applications, 2003, Proceedings, 14th Internation 1-5 Sept. 2003 Page(s):536 - 541 Digital Object Identifier 10.1109/DEXA.2003.1232078 AbstractPlus | Full Text: PDF(244 KB) IEEE CNF Rights and Permissions 7. An update-risk based approach to TTL estimation in Web caching Jeong-Joon Lee; Kyu-Young Whang; Byung Suk Lee; Ji-Woong Chang; Web Information Systems Engineering, 2002, WISE 2002, Proceedings of the International Conference on 12-14 Dec. 2002 Page(s):21 - 29 Digital Object Identifier 10.1109/WISE.2002.1181640 AbstractPlus | Full Text: PDF(358 KB) IEEE CNF Rights and Permissions 8. Internet user behavior: compared study of the access traces and applical discovery of communities Lancieri, L.; Durand, N.; Systems, Man and Cybernetics, Part A, IEEE Transactions on Volume 36, Issue 1, Jan. 2006 Page(s):208 - 219 Digital Object Identifier 10.1109/TSMCA.2005.859095 AbstractPlus | Full Text: PDF(312 KB) IEEE JNL Rights and Permissions 9. Web services: e-commerce partner integration Kulkarni, N.; Kumar, S.; Mani, K.; Padmanabhuni, S.; IT Professional Volume 7, Issue 2, Mar-Apr 2005 Page(s):23 - 28 Digital Object Identifier 10.1109/MITP.2005.46 AbstractPlus | Full Text: PDF(8008 KB) | IEEE JNL Rights and Permissions 10. Proxy cache algorithms: design, implementation, and performance Shim, J.; Scheuermann, P.; Vingralek, R.; Knowledge and Data Engineering, IEEE Transactions on Volume 11, Issue 4, July-Aug. 1999 Page(s):549 - 562 Digital Object Identifier 10.1109/69.790804 AbstractPlus | References | Full Text: PDF(1132 KB) IEEE JNL Rights and Permissions 11. Improving proxy cache performance: analysis of three replacement polic Dilley, J.; Arlitt, M.; Internet Computing, IEEE Volume 3, Issue 6, Nov.-Dec. 1999 Page(s):44 - 50 Digital Object Identifier 10.1109/4236.806998 AbstractPlus | References | Full Text: PDF(436 KB) | IEEE JNL Rights and Permissions 12. Delayed Internet routing convergence Labovitz, C.; Ahuja, A.; Bose, A.; Jahanian, F.; Networking, IEEE/ACM Transactions on

Volume 9, Issue 3, June 2001 Page(s):293 - 306 Digital Object Identifier 10.1109/90.929852

	AbstractPlus   References   Full Text: PDF(220 KB) 기본문문 JNL Rights and Permissions
	13. Wise drives [hard disk drive] Hughes, G.F.; Spectrum. IEEE Volume 39, Issue 8, Aug. 2002 Page(s):37 - 41 Digital Object Identifier 10.1109/MSPEC.2002.1021942
	AbstractPlus   Full Text: <u>PDF(310 KB)   Full Text: HTML</u> IEEE JNI. Rights and Permissions
	14. On the geographic location of Internet resources Lakhina, A.; Byers, J.W.; Crovella, M.; Matta, I.; Selected Areas in Communications. IEEE Journal on Volume 21, Issue 6, Aug. 2003 Page(s):934 - 948 Digital Object Identifier 10.1109/JSAC.2003.814667
	AbstractPlus   References   Full Text: PDE(981 KB)   IEEE JNL Rights and Permissions
	15. File system caching in large point-to-point networks Austin, P.B.; Murray, K.A.; Wellings, A.J.; Software Engineering Journal Volume 7, Issue 1, Jan. 1992 Page(s):65 - 80
·	AbstractPlus   Full Text: PDF(1048 KB) IEE JNL
	16. A distributed algorithm for sharing Web cache disk capacity Alyfantis, G.; Hadjiefthymiades, S.; Merakos, L.; Kostopoulos, P.; Parallel and Distributed Systems, 2006, ICPADS 2006, 12th International Conty Volume 1, 12-15 July 2006 Page(s):8 pp. Digital Object Identifier 10.1109/ICPADS.2006.12
	AbstractPlus   Full Text: PDF(304 KB) IEEE CNF Rights and Permissions
	17. A cross-architectural interface for code cache manipulation Hazelwood, K.; Cohn, R.; Code Generation and Optimization, 2006, CGO 2006, International Symposium 26-29 March 2006 Page(s):11 pp. Digital Object Identifier 10.1109/CGO.2006.3  AbstractPlus   Full Text: PDF(408 KB)   IEEE CNF
	Rights and Permissions
	18. Utilizing layered multicast for Web caching Schmidbauer, J.; Linder, L.; Web Content Caching and Distribution, 2005, WCW 2005, 10th International V 12-13 Sept. 2005 Page(s):121 - 132 Digital Object Identifier 10.1109/WCW.2005.22
	AbstractPlus   Full Text: PDF(752 KB) IEEE CNF Rights and Permissions
	19. Emulating Web cache replacement algorithms versus a real system Cardenas, L.G.; Gil, J.A.; Sahuquillo, J.; Pont, A.; Computers and Communications, 2005. ISCC 2005. Proceedings, 10th IEEE.\$ 27-30 June 2005 Page(s):891 - 897 Digital Object Identifier 10.1109/ISCC.2005.63
	AbstractPlus   Full Text: PDF(464 KB) 1888 CNF Rights and Permissions
	20.

Clarke, D.; Suh, G.E.; Gassend, B.; Sudan, A.; van Dijk, M.; Devadas, S.; Security and Privacy, 2005 IEEE Symposium on 8-11 May 2005 Page(s):139 - 153 Digital Object Identifier 10.1109/SP.2005.24 AbstractPlus | Full Text: PDF(232 KB) IEEE CNF Rights and Permissions 21. Client clustering for traffic and location estimation Amini, L.; Schulzrinne, H.; Distributed Computing Systems, 2004, Proceedings, 24th International Confer-2004 Page(s):730 - 737 Digital Object Identifier 10.1109/ICDCS.2004.1281641 AbstractPlus | Full Text: PDF(658 KB) IEEE CNF Rights and Permissions 22. Integrated network traffic measurement and billing system Yong, K.C.; Lai, Z.S.; Ananda, A.L.; Networks, 2003, ICON2003. The 11th IEEE International Conference on 28 Sept.-1 Oct. 2003 Page(s):19 - 24 Digital Object Identifier 10.1109/ICON.2003.1266161 AbstractPlus | Full Text: PDF(470 KB) IEEE CNF Rights and Permissions 23. Building a flexile Web caching system Sosa, V.J.; Gonzalez, G.; Navarro, L.; Computer Science, 2003, ENC 2003, Proceedings of the Fourth Mexican Inter Conference on 8-12 Sept. 2003 Page(s):60 - 65 Digital Object Identifier 10.1109/ENC.2003.1232875 AbstractPlus | Full Text: PDF(328 KB) IEEE CNF Rights and Permissions 24. A log-based write-back mechanism for cooperative caching Narasimhan, S.; Sohoni, S.; Yiming Hu; Parallel and Distributed Processing Symposium, 2003, Proceedings, Internation 22-26 April 2003 Page(s):10 pp. Digital Object Identifier 10.1109/IPDPS.2003.1213157 AbstractPlus | Full Text: PDF(296 KB) IEEE CNF Rights and Permissions 25. DMASC: a tool for visualizing user paths through a Web site Stones, C.; Sobol, S.; Database and Expert Systems Applications, 2002, Proceedings, 13th Internative 2-6 Sept. 2002 Page(s):389 - 393 AbstractPlus | Full Text: PDF(650 KB) IEEE CNF Rights and Permissions

inspec"

Help Contact Us Privacy &:

© Copyright 2006 IEEE -



Home | Login | Logout | Access Information | Alerts |

#### Welcome United States Patent and Trademark Office

Search Results

BROWSE

SEARCH

IEEE XPLORE GUIDE

Results for "(cache log, discare	d queue <and>refresh queue)"</and>
----------------------------------	------------------------------------

Your search matched 28 of 1484991 documents.

A maximum of 28 results are displayed, 25 to a page, sorted by Relevance in Descending order.

**Modify Search** 

» Search Options

View Session History

New Search

(cache log, discard queue<and>refresh queue)

Check to search only within this results set

Search

☑ e-mail

» Key

IEEE Journal or

Magazine

IEE JNL

IEEE JNL

IEE Journal or Magazine

HEER CNF

IEEE Conference

roccounty

iee Cnf

IEE Conference Proceeding

IEEE STD IEEE Standard

view selected items

Select All Deselect All

26. Performance of Web proxy caching in heterogeneous bandwidth environ

26. Performance of Web proxy caching in heterogeneous bandwick Feldmann, A.; Caceres, R.; Douglis, F.; Glass, G.; Rabinovich, M.;

INFOCOM '99. Eighteenth Annual Joint Conference of the IEEE Computer and

Societies. Proceedings, IEEE

Volume 1, 21-25 March 1999 Page(s):107 - 116 vol.1 Digital Object Identifier 10.1109/INFCOM.1999.749258

AbstractPlus | Full Text: PDF(996 KB) IEEE CNF

Rights and Permissions

27. Log-structured file systems

Douglis, F.; Ousterhout, J.;

COMPCON Spring '89. Thirty-Fourth IEEE Computer Society International Controllectual Leverage, Digest of Papers.

27 Feb.-3 March 1989 Page(s):124 - 129

Digital Object Identifier 10.1109/CMPCON.1989.301914

AbstractPlus | Full Text: PDF(504 KB) IEEE CNF

Rights and Permissions

28. Cache Logic FPGAs for building adaptive hardware

Camarota, R.; Rosenberg, J.;

Field Programmable Gate Arrays - Technology and Applications, IEE Colloquia

15 Feb 1993 Page(s):1/1 - 1/3

AbstractPlus | Full Text: PDF(144 KB) IEE CNF

Help Contact Us Privacy &:

@ Copyright 2006 IEEE -

™owby Ѿinspec"



Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

+discard +queue, +cache +miss, +refresh +queue, +data +of



### THE ACM DIGITAL LIBRARY

Feedback Report a problem Salisfaction survey

Published before October 2003

Found

Terms used

**12** o

discard queue cache miss refresh queue data object cache log prefetch predict preload lookahead 146,827

Sort results by

relevance 👻

Save results to a Binder

Search Tips

Try an Advanced Search

Try this search in <u>The ACM Guide</u>

Display results

expanded form

Open results in a new window

Results 1 - 12 of 12

Relevance scale

An architecture to support scalable online personalization on the Web
Anindya Datta, Kaushik Dutta, Debra VanderMeer, Krithi Ramamritham, Shamkant B. Navathe
August 2001 The VLDB Journal — The International Journal on Very Large Data Bases,
Volume 10 Issue 1

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(167.25 KB) Additional Information: full citation, abstract, citings, index terms

Online personalization is of great interest to e-companies. Virtually all personalization technologies are based on the idea of storing as much historical customer session data as possible, and then querying the data store as customers navigate through a web site. The holy grail of online personalization is an environment where fine-grained, detailed historical session data can be queried based on current online navigation patterns for use in formulating real-time responses. Unfortunately, as mo ...

**Keywords:** Behavior-based personalization, Dynamic lookahead profile, Profile caching, Scalable online personalization, Web site and interaction model

<sup>2</sup> Enabling scalable online personalization on the Web

Debra VanderMeer, Kaushik Dutta, Anindya Datta, Krithi Ramamritham, Shamkant B. Navanthe

October 2000 Proceedings of the 2nd ACM conference on Electronic commerce EC '00

**Publisher: ACM Press** 

Full text available: pdf(491,69 KB) Additional Information: full citation, references, citings, index terms

Keywords: dynamic profiling, e-commerce, online personalization, user behavior

3 Engineering web cache consistency

Jian Yin, Lorenzo Alvisi, Mike Dahlin, Arun Iyengar

August 2002 ACM Transactions on Internet Technology (TOIT), Volume 2 Issue 3

Publisher: ACM Press

Full text available: pdf(403.96 KB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

Server-driven consistency protocols can reduce read latency and improve data freshness

for a given network and server overhead, compared to the traditional consistency protocols that rely on client polling. Server-driven consistency protocols appear particularly attractive for large-scale dynamic Web workloads because dynamically generated data can change rapidly and unpredictably. However, there have been few reports on engineering server-driven consistency for such workloads. This article repo ...

**Keywords**: Cache coherence, cache consistency, dynamic content, lease, scalability, volume

4 GPRSWeb: optimizing the web for GPRS links

Rajiv Chakravorty, Andrew Clark, Ian Pratt

May 2003 Proceedings of the 1st international conference on Mobile systems, applications and services MobiSys '03

**Publisher: ACM Press** 

Full text available: pdf(1.03 MB) Additional Information: full citation, abstract, references, citings

The General Packet Radio Service (GPRS) is being deployed by GSM network operators world-wide, and promises to offer users "always-on" data access at bandwidths comparable to that of conventional fixed-line telephone modems. Unfortunately, many users have found the reality to be rather different, experiencing very disappointing performance when, for example, browsing the web over GPRS.In this paper we investigate what causes the HTTP protocol and its underlying transport TCP to underperform in a ...

5 End-to-end WAN service availability

Michael Dahlin, Bharat Baddepudi V. Chandra, Lei Gao, Amol Nayate April 2003 IEEE/ACM Transactions on Networking (TON), Volume 11 Issue 2

Publisher: IEEE Press

Full text available: pdf(684.07 KB)

Additional Information: full cliation, abstract, references, citings, index terms

This paper seeks to understand how network failures affect the availability of service delivery across wide-area networks (WANs) and to evaluate classes of techniques for improving end-to-end service availability. Using several large-scale connectivity traces, we develop a model of network unavailability that includes key parameters such as failure location and failure duration. We then use trace-based simulation to evaluate several classes of techniques for coping with network unavailability. W ...

**Keywords:** availability, disconnected operation, failure model, internet, overlay routing, replication, world-wide web

<sup>6</sup> The SimpleScalar tool set, version 2.0

Doug Burger, Todd M. Austin

June 1997 ACM SIGARCH Computer Architecture News, Volume 25 Issue 3

**Publisher: ACM Press** 

Full text available: pdf(985.46 KB) Additional Information: full citation, abstract, citings, index terms

This document describes release 2.0 of the SimpleScalar tool set, a suite of free, publicly available simulation tools that offer both detailed and high-performance simulation of modern microprocessors. The new release offers more tools and capabilities, precompiled binaries, cleaner interfaces, better documentation, easier installation, improved portability, and higher performance. This paper contains a complete description of the tool set, including retrieval and installation instructions, a d ...

7 Practical byzantine fault tolerance and proactive recovery

Miguel Castro, Barbara Liskov

November 2002 ACM Transactions on Computer Systems (TOCS), Volume 20 Issue 4

Publisher: ACM Press

Full text available: pdf(1.63 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> <u>terms</u>, <u>review</u>

Our growing reliance on online services accessible on the Internet demands highly available systems that provide correct service without interruptions. Software bugs, operator mistakes, and malicious attacks are a major cause of service interruptions and they can cause arbitrary behavior, that is, Byzantine faults. This article describes a new replication algorithm, BFT, that can be used to build highly available systems that tolerate Byzantine faults. BFT can be used in practice to implement re ...

**Keywords**: Byzantine fault tolerance, asynchronous systems, proactive recovery, state machine replication, state transfer

8 System-level power optimization: techniques and tools

Luca Benini, Giovanni de Micheli

April 2000 ACM Transactions on Design Automation of Electronic Systems (TODAES),

Volume 5 Issue 2
Publisher: ACM Press

Full text available: mpdf(385.22 KB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

This tutorial surveys design methods for energy-efficient system-level design. We consider electronic sytems consisting of a hardware platform and software layers. We consider the three major constituents of hardware that consume energy, namely computation, communication, and storage units, and we review methods of reducing their energy consumption. We also study models for analyzing the energy cost of software, and methods for energy-efficient software design and compilation. This survery ...

9 Update propagation strategies to improve freshness in lazy master replicated databases

Esther Pacitti, Eric Simon

February 2000 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 8 Issue 3-4

Publisher: Springer-Verlag New York, Inc.

Full text available: 📆 pdf(151,35 KB) Additional Information: full citation, abstract, citings, index terms

Many distributed database applications need to replicate data to improve data availability and query response time. The two-phase commit protocol guarantees mutual consistency of replicated data but does not provide good performance. Lazy replication has been used as an alternative solution in several types of applications such as on-line financial transactions and telecommunication systems. In this case, mutual consistency is relaxed and the concept of freshness is used to measure the deviation ...

Keywords: Data replication, Distributed databases, Performance evaluation

10 Higher-order distributed objects

Henry Cejtin, Suresh Jagannathan, Richard Kelsey

September 1995 ACM Transactions on Programming Languages and Systems (TOPLAS), Volume 17 Issue 5

Publisher: ACM Press

Full text available: Mpdf(2.33 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms, <u>review</u>

http://portal.acm.org/results.cfm?CFID=13394798&CFTOKEN=93818233&adv=1&COLL... 2/2/2007

We describe a distributed implementation of Scheme that permits efficient transmission of higher-order objects such as closures and continuations. The integration of distributed communication facilities within a higher-order programming language engenders a number of new abstractions and paradigms for distributed computing. Among these are userspecified load-balancing and migration policies for threads, incrementally linked distributed computations, and parameterized client-server applicat ...

Keywords: concurrency, continuations, higher-order languages, message-passing

11 Layered transmission and caching for the multicast session directory service

Andrew Swan, Steven McCanne, Lawrence A. Rowe

September 1998 Proceedings of the sixth ACM international conference on Multimedia **MULTIMEDIA '98** 

Publisher: ACM Press

Full text available: pdf(1.32 MB) Additional Information: full citation, references, citings, index terms

12 A potpourri of ideas for event-based processing: A case study on event dissemination

in an active overlay network environment Sérgio Duarte, J. Legatheaux Martins, Henrique J. Domingos, Nuno Preguiça

June 2003 Proceedings of the 2nd international workshop on Distributed event-based systems DEBS '03

Publisher: ACM Press

Full text available: pdf(1.60 MB) Additional Information: full citation, abstract, references

In this paper, we describe a case study of the design and development of a groupconferencing tool suite, built on top of an overlay network based event dissemination framework, which is extensible via quality of service template plug-ins. We explain, for each of the tools, how the framework built-in conveniences were explored to create simple but effective distributed solutions, backed by the appropriate quality of service templates, whose design we also discuss.

Keywords: active networks, case study, event dissemination, multicasting, overlay networks, quality of service (QoS)

Results 1 - 12 of 12

· The ACM Portal is published by the Association for Computing Machinery: Copyright © 2007 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player

Desario	Subscribe (Full Service) Register (Limited Service, Free) Login	
● P®RTAL	Search: © The ACM Digital Library C The Guide	
USPTO		
THE ACM DIGITAL LIBRARY	Advanced Search Tips	
Enter words, phrases or names below. Surround phrases or full names with double quotation marks.		
Desired Results: must have all of the words or phrases discard, cache miss, refresh, data object, cache must have any of the words or phrases prefetch, predict, preload, lookahead must have none of the words or phrases	Name or Affiliation:  Authored by: all Cany Cnone  Edited by: all Cany Cnone  Reviewed by: all Cany Cnone	
Only search in:*  O Title O Abstract O Review O All	nformation	
	e information, including full text where available, unless specified	
ISBN / ISSN: ● Exact ○ Expand	DOI: € Exact C Expand	
Published:	Conference Proceeding:	
By:   all Cany Cnone	Sponsored By:	
In: ® all C any C none	Conference Location:	
Since:   Month   Year   Period   Period	Conference Year:	
As: Any type of publication		
Classification: (CCS) Primary Only	Results must have accessible:	
Classified as: ® all □ any □ none	Full Text Abstract Review	
Subject Descriptor:   all  any  none	······································	
Keyword Assigned: © all C any C none		



Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

+discard, +cache +miss, +refresh, +data +object, +cache +lc



### THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Published before October 2003

Terms used

Found 30 of 146,827

discard cache miss refresh data object cache log prefetch predict preload lookahead

Sort results bγ

relevance

Save results to a Binder Search Tips

Try an Advanced Search Try this search in The ACM Guide

Display results

expanded form

Open results in a new window

Results 1 - 20 of 30

Result page: 1 2

next

Relevance scale 

Relevance

An architecture to support scalable online personalization on the Web Anindya Datta, Kaushik Dutta, Debra VanderMeer, Krithi Ramamritham, Shamkant B. Navathe

August 2001 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 10 Issue 1

Publisher: Springer-Verlag New York, Inc.

Full text available: 📆 pdi(167.25 KB) Additional Information: full citation, abstract, citings, index terms

Online personalization is of great interest to e-companies. Virtually all personalization technologies are based on the idea of storing as much historical customer session data as possible, and then querying the data store as customers navigate through a web site. The holy grail of online personalization is an environment where fine-grained, detailed historical session data can be queried based on current online navigation patterns for use in formulating real-time responses. Unfortunately, as mo ...

Keywords: Behavior-based personalization, Dynamic lookahead profile, Profile caching, Scalable online personalization, Web site and interaction model

<sup>2</sup> Enabling scalable online personalization on the Web



Debra VanderMeer, Kaushik Dutta, Anindya Datta, Krithi Ramamritham, Shamkant B. Navanthe

October 2000 Proceedings of the 2nd ACM conference on Electronic commerce EC '00

Publisher: ACM Press

Full text available: ndf(491.69 KB) Additional Information: full citation, references, citings, index terms

Keywords: dynamic profiling, e-commerce, online personalization, user behavior

Engineering web cache consistency

Jian Yin, Lorenzo Alvisi, Mike Dahlin, Arun Iyengar

August 2002 ACM Transactions on Internet Technology (TOIT), Volume 2 Issue 3

Publisher: ACM Press

Additional Information: full citation, abstract, references, citings, index

Full text available: mpdf(403.96 KB)

Server-driven consistency protocols can reduce read latency and improve data freshness for a given network and server overhead, compared to the traditional consistency protocols that rely on client polling. Server-driven consistency protocols appear particularly attractive for large-scale dynamic Web workloads because dynamically generated data can change rapidly and unpredictably. However, there have been few reports on engineering server-driven consistency for such workloads. This article repo ...

Keywords: Cache coherence, cache consistency, dynamic content, lease, scalability, volume

Aging through cascaded caches: performance issues in the distribution of web



content

Edith Cohen, Haim Kaplan

August 2001 ACM SIGCOMM Computer Communication Review, Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications SIGCOMM '01, Volume 31 Issue 4

Publisher: ACM Press

Full text available: pdf(327.13 KB)

Additional Information: full cilation, abstract, references, citings, index terms

The Web is a distributed system, where data is stored and disseminated from both origin servers and caches. Origin servers provide the most up-to-date copy whereas caches store and serve copies that had been cached for a while. Origin servers do not maintain per-client state, and weak-consistency of cached copies is maintained by the origin server attaching to each copy an expiration time. Typically, the lifetime-duration of an object is fixed, and as a result, a copy fetched direc ...

Decentralizing a global naming service for improved performance and fault tolerance D. R. Cheriton, T. P. Mann



May 1989 ACM Transactions on Computer Systems (TOCS), Volume 7 Issue 2

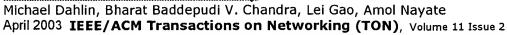
Publisher: ACM Press

Full text available: pdf(3.19 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

Naming is an important aspect of distributed system design. A naming system allows users and programs to assign character-string names to objects, and subsequently use the names to refer to those objects. With the interconnection of clusters of computers by wide-area networks and internetworks, the domain over which naming systems must function is growing to encompass the entire world. In this paper we address the problem of a global naming system, proposing a three-level naming ...

End-to-end WAN service availability



Publisher: IEEE Press

Full text available: pdf(684.07 KB)

Additional Information: full citation, abstract, references, citings, index terms

This paper seeks to understand how network failures affect the availability of service delivery across wide-area networks (WANs) and to evaluate classes of techniques for improving end-to-end service availability. Using several large-scale connectivity traces, we develop a model of network unavailability that includes key parameters such as failure location and failure duration. We then use trace-based simulation to evaluate several classes of techniques for coping with network unavailability. W ...

**Keywords**: availability, disconnected operation, failure model, internet, overlay routing, replication, world-wide web

7 GPRSWeb: optimizing the web for GPRS links

Rajiv Chakravorty, Andrew Clark, Ian Pratt

May 2003 Proceedings of the 1st international conference on Mobile systems, applications and services MobiSys '03

Publisher: ACM Press

Full text available: pdf(1.03 MB) Additional Information: full citation, abstract, references, citings

The General Packet Radio Service (GPRS) is being deployed by GSM network operators world-wide, and promises to offer users "always-on" data access at bandwidths comparable to that of conventional fixed-line telephone modems. Unfortunately, many users have found the reality to be rather different, experiencing very disappointing performance when, for example, browsing the web over GPRS. In this paper we investigate what causes the HTTP protocol and its underlying transport TCP to underperform in a ...

Information and control in gray-box systems

Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau

October 2001 ACM SIGOPS Operating Systems Review, Proceedings of the eighteenth ACM symposium on Operating systems principles SOSP '01, Volume 35 Issue

Publisher: ACM Press

Full text available: pdf(1.59 MB)

Additional Information: full citation, abstract, references, citings, index

In modern systems, developers are often unable to modify the underlying operating system. To build services in such an environment, we advocate the use of gray-box techniques. When treating the operating system as a gray-box, one recognizes that not changing the OS restricts, but does not completely obviate, both the information one can acquire about the internal state of the OS and the control one can impose on the OS. In this paper, we develop and investigate three gray-bo ...

A predicate-based caching scheme for client-server database architectures Arthur M. Keller, Julie Basu

January 1996 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 5 Issue 1

Publisher: Springer-Verlag New York, Inc.

Full text available: ndf(162.80 KB) Additional Information: full citation, abstract, citings, index terms

We propose a new client-side data-caching scheme for relational databases with a central server and multiple clients. Data are loaded into each client cache based on queries executed on the central database at the server. These queries are used to form predicates that describe the cache contents. A subsequent query at the client may be satisfied in its local cache if we can determine that the query result is entirely contained in the cache. This issue is called *cache completeness*. A separ ...

Keywords: Cache completeness, Cache currency, Caching, Multiple clients, Relational databases

10 Organization of invalidation reports for energy-efficient cache invalidation in mobile environments Kian-Lee Tan



June 2001 Mobile Networks and Applications, Volume 6 Issue 3

Publisher: Kluwer Academic Publishers

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(190.83 KB) terms

In a wireless environment, mobile clients often cache frequently accessed data to reduce contention on the limited wireless bandwidth. However, it is difficult for clients to ascertain the validity of their cache content because of their frequent disconnection. One promising cache invalidation approach is the Bit-Sequences scheme that organizes invalidation reports as a set of binary bit sequences with an associated set of timestamps. The report is periodically broadcast by ...

Keywords: access time, bit-sequences, cache invalidation, disconnection, energy consumption, mobile computing

11 System-level power optimization: techniques and tools

Luca Benini, Giovanni de Micheli

April 2000 ACM Transactions on Design Automation of Electronic Systems (TODAES), Volume 5 Issue 2

**Publisher: ACM Press** 

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(385,22 KB)

This tutorial surveys design methods for energy-efficient system-level design. We consider electronic sytems consisting of a hardware platform and software layers. We consider the three major constituents of hardware that consume energy, namely computation, communication, and storage units, and we review methods of reducing their energy consumption. We also study models for analyzing the energy cost of software, and methods for energy-efficient software design and compilation. This survery ...

12 Authentication in distributed systems: theory and practice

Butler Lampson, Martín Abadi, Michael Burrows, Edward Wobber

November 1992 ACM Transactions on Computer Systems (TOCS), Volume 10 Issue 4

Publisher: ACM Press

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(3.37 MB) terms, review

We describe a theory of authentication and a system that implements it. Our theory is based on the notion of principal and a "speaks for" relation between principals. A simple principal either has a name or is a communication channel; a compound principal can express an adopted role or delegated authority. The theory shows how to reason about a principal's authority by deducing the other principals that it can speak for; authenticating a channel is one important application. We ...

Keywords: certification authority, delegation, group, interprocess communication, key distribution, loading programs, path name, principal, role, secure channel, speaks for, trusted computing base

Intelligent database caching through the use of page-answers and page-traces Nabil Kamel, Roger King

December 1992 ACM Transactions on Database Systems (TODS), Volume 17 Issue 4

Publisher: ACM Press

Additional Information: full citation, abstract, references, citings, index Full text available: R pdf(3.08 MB) terms

In this paper a new method to improve the utilization of main memory systems is

presented. The new method is based on prestoring in main memory a number of query answers, each evaluated out of a single memory page. To this end, the ideas of pageanswers and page-traces are formally described and their properties analyzed. The query model used here allows for selection, projection, join, recursive queries as well as arbitrary combinations. We also show how to apply the approach under update ...

**Keywords**: artificial intelligence, databases, page access

14 Practical byzantine fault tolerance and proactive recovery

Miguel Castro, Barbara Liskov
November 2002 ACM Transactions on Computer Systems (TOCS), Volume 20 Issue 4

Publisher: ACM Press

Full text available: pdf(1.63 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

Our growing reliance on online services accessible on the Internet demands highly available systems that provide correct service without interruptions. Software bugs, operator mistakes, and malicious attacks are a major cause of service interruptions and they can cause arbitrary behavior, that is, Byzantine faults. This article describes a new replication algorithm, BFT, that can be used to build highly available systems that tolerate Byzantine faults. BFT can be used in practice to implement re ...

Keywords: Byzantine fault tolerance, asynchronous systems, proactive recovery, state machine replication, state transfer

15 The cache location problem

P. Krishnan, Danny Raz, Yuval Shavitt

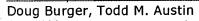
October 2000 IEEE/ACM Transactions on Networking (TON), Volume 8 Issue 5

Publisher: IEEE Press

Full text available: pdf(470.11 KB) Additional Information: full citation, references, citings, index terms

Keywords: location problem, mirror placement, transparent cache

16 The SimpleScalar tool set, version 2.0



June 1997 ACM SIGARCH Computer Architecture News, Volume 25 Issue 3

**Publisher: ACM Press** 

Full text available: pdf(985,46 KB) Additional Information: full citation, abstract, citings, index terms

This document describes release 2.0 of the SimpleScalar tool set, a suite of free, publicly available simulation tools that offer both detailed and high-performance simulation of modern microprocessors. The new release offers more tools and capabilities, precompiled binaries, cleaner interfaces, better documentation, easier installation, improved portability, and higher performance. This paper contains a complete description of the tool set, including retrieval and installation instructions, a d ...

17 Report on the fourth ACM SIGOPS European workshop fault tolerance support in





distributed systems Özalp Babaoğlu

January 1991 ACM SIGOPS Operating Systems Review, Volume 25 Issue 1

Publisher: ACM Press

Additional Information: full citation, index ferms Full text available: pdf(1.76 MB)

18 Enabling dynamic content caching for database-driven web sites

K. Selçuk Candan, Wen-Syan Li, Qiong Luo, Wang-Pin Hsiung, Divyakant Agrawal

May 2001 ACM SIGMOD Record, Proceedings of the 2001 ACM SIGMOD international conference on Management of data SIGMOD '01, Volume 30 Issue 2

**Publisher: ACM Press** 

Full text available: pdf(319.67 KB)

Additional Information: full citation, abstract, references, citings, index terms

Web performance is a key differentiation among content providers. Snafus and slowdowns at major web sites demonstrate the difficulty that companies face trying to scale to a large amount of web traffic. One solution to this problem is to store web content at server-side and edge-caches for fast delivery to the end users. However, for many ecommerce sites, web pages are created dynamically based on the current state of business processes, represented in application servers and databases

Keywords: JDBC, application server, database driven web site, dynamic content caching, invalidation, web acceleration

The X window system

Robert W. Scheifler, Jim Gettys

April 1986 ACM Transactions on Graphics (TOG), Volume 5 Issue 2

Publisher: ACM Press

Full text available: pdf(2.76 MB)

Additional Information: full citation, abstract, references, citings, index

terms, review

An overview of the X Window System is presented, focusing on the system substrate and the low-level facilities provided to build applications and to manage the desktop. The system provides high-performance, high-level, device-independent graphics. A hierarchy of resizable, overlapping windows allows a wide variety of application and user interfaces to be built easily. Network-transparent access to the display provides an important degree of functional separation, without significantly affec ...

20 Log-based receiver-reliable multicast for distributed interactive simulation

Hugh W. Holbrook, Sandeep K. Singhal, David R. Cheriton

October 1995 ACM SIGCOMM Computer Communication Review , Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication SIGCOMM '95, Volume 25 Issue 4

Publisher: ACM Press

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(1.44 MB)

Reliable multicast communication is important in large-scale distributed applications. For example, reliable multicast is used to transmit terrain and environmental updates in distributed simulations. To date, proposed protocols have not supported these applications' requirements, which include wide-area data distribution, low-latency packet loss detection and recovery, and minimal data and management over-head within finegrained multicast groups, each containing a single data source. In this pa ...

Results 1 - 20 of 30 Result page: 1 2 next

> The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player



Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

+discard, +cache +miss, +refresh, +data +object, +cache +lc



### THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Published before October 2003

Terms used

Found 30 of 146,827

<u>discard cache miss refresh data object cache log prefetch predict preload lookahead</u>

Sort results by

Display

relevance

expanded form

Save results to a Binder 2 Search Tips

Try an Advanced Search Try this search in The ACM Guide

results .

Open results in a new window

Results 21 - 30 of 30

Result page: previous 1 2

21 A case for dynamic view management

Yannis Kotidis, Nick Roussopoulos

December 2001 ACM Transactions on Database Systems (TODS), Volume 26 Issue 4

**Publisher: ACM Press** 

Full text available: pdf(892.57 K8)

Additional Information: full citation, abstract, references, citings, index terms, review

Materialized aggregate views represent a set of redundant entities in a data warehouse that are frequently used to accelerate On-Line Analytical Processing (OLAP). Due to the complex structure of the data warehouse and the different profiles of the users who submit queries, there is need for tools that will automate and ease the view selection and management processes. In this article we present DynaMat, a system that manages dynamic collections of materialized aggregate views in a data warehous ...

**Keywords**: Data cube, OLAP, data warehousing, materialized views

22 Towards effective and efficient free space management

Mark L. McAuliffe, Michael J. Carey, Marvin H. Solomon

June 1996 ACM SIGMOD Record, Proceedings of the 1996 ACM SIGMOD international conference on Management of data SIGMOD '96. Volume 25 Issue 2

**Publisher: ACM Press** 

Full text available: pdf(1.34 MB)

Additional Information: full citation, abstract, references, citings, index <u>terms</u>

An important problem faced by many database management systems is the "online object placement problem"--the problem of choosing a disk page to hold a newly allocated object. In the absence of clustering criteria, the goal is to maximize storage utilization. For main-memory based systems, simple heuristics exist that provide reasonable space utilization in the worst case and excellent utilization in typical cases. However, the storage management problem for databases includes significant additio ...

23 Update propagation strategies to improve freshness in lazy master replicated databases

Esther Pacitti, Eric Simon

February 2000 The VLDB Journal — The International Journal on Very Large Data

Bases, Volume 8 Issue 3-4

Publisher: Springer-Verlag New York, Inc.

Full text available: ndf(151.35 KB) Additional Information: full citation, abstract, citings, index terms

Many distributed database applications need to replicate data to improve data availability and query response time. The two-phase commit protocol guarantees mutual consistency of replicated data but does not provide good performance. Lazy replication has been used as an alternative solution in several types of applications such as on-line financial transactions and telecommunication systems. In this case, mutual consistency is relaxed and the concept of freshness is used to measure the deviation ...

Keywords: Data replication, Distributed databases, Performance evaluation

24 Layered transmission and caching for the multicast session directory service

Andrew Swan, Steven McCanne, Lawrence A. Rowe
September 1998 Proceedings of the sixth ACM international conference on Multimedia
MULTIMEDIA '98

**Publisher: ACM Press** 

Full text available: pdf(1.32 MB)

Additional Information: full citation, references, citings, index terms

25 <u>Virtual memory management for database systems</u>

irving L. Traiger

October 1982 ACM SIGOPS Operating Systems Review, Volume 16 Issue 4

**Publisher: ACM Press** 

Full text available: pdf(2.08 MB) Additional Information: full citation, abstract, references, citings

Over the last several years, a number of hardware and software systems have been developed which map entire files directly into the virtual memory address spaces used by programs. Since all file contents are directly addressable, there is no need for a programmer to issue explicit file system actions, such as Read or Write. In addition, all of the buffer management problems are eliminated, since programmers do not have to squeeze pieces of large files into small virtual spaces. Although these ad ...

26 Authentication in the Taos operating system

Edward Wobber, Martín Abadi, Michael Burrows, Butler Lampson
February 1994 ACM Transactions on Computer Systems (TOCS), Volume 12 Issue 1

Publisher: ACM Press

Full text available: pdf(1.88 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

We describe a design for security in a distributed system and its implementation. In our design, applications gain access to security services through a narrow interface. This interface provides a notion of identity that includes simple principals, groups, roles, and delegations. A new operating system component manages principals, credentials, and secure channels. It checks credentials according to the formal rules of a logic of authentication. Our implementation is efficient enough to sup ...

Keywords: cryptography, mathematical logic

27 Higher-order distributed objects

Henry Cejtin, Suresh Jagannathan, Richard Kelsey
September 1995 ACM Transactions on Programming Languages and Systems

(TOPLAS), Volume 17 Issue 5

Results (page 2): +discard, +cache +miss, +refresh, +data +object, +cache +log prefetch, ... Page 3 of 4

Publisher: ACM Press

Full text available: mpdf(2.33 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms, review

We describe a distributed implementation of Scheme that permits efficient transmission of higher-order objects such as closures and continuations. The integration of distributed communication facilities within a higher-order programming language engenders a number of new abstractions and paradigms for distributed computing. Among these are user-specified load-balancing and migration policies for threads, incrementally linked distributed computations, and parameterized client-server applicat ...

Keywords: concurrency, continuations, higher-order languages, message-passing

28 Ad hoc and sensor networks: Distributed algorithms for guiding navigation across a



sensor network

Qun Li, Michael De Rosa, Daniela Rus

September 2003 Proceedings of the 9th annual international conference on Mobile computing and networking MobiCom '03

Publisher: ACM Press.

Full text available: pdf(810.66 KB)

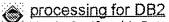
Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> <u>terms</u>

We develop distributed algorithms for self-organizing sensor networks that respond to directing a target through a region. The sensor network models the danger levels sensed across its area and has the ability to adapt to changes. It represents the dangerous areas as obstacles. A protocol that combines the artificial potential field of the sensors with the goal location for the moving object guides the object incrementally across the network to the goal, while maintaining the safest distance to ...

Keywords: motes, navigation, potential field, robotics, sensor networks

29 Industrial sessions: beyond relational tables: Garlic: a new flavor of federated query





Vanja Josifovski, Peter Schwarz, Laura Haas, Eileen Lin

June 2002 Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02

Publisher: ACM Press

Full text available: pdf(1.05 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

In a large modern enterprise, information is almost inevitably distributed among several database management systems. Despite considerable attention from the research community, relatively few commercial systems have attempted to address this issue. This paper describes new technology that enables clients of IBM's DB2 Universal Database to access the data and specialized computational capabilities of a wide range of non-relational data sources. This technology, based on the Garlic prototype deve ...

30 A potpourri of ideas for event-based processing: A case study on event dissemination



in an active overlay network environment

Sérgio Duarte, J. Legatheaux Martins, Henrique J. Domingos, Nuno Preguiça

June 2003 Proceedings of the 2nd international workshop on Distributed eventbased systems DEBS '03

**Publisher: ACM Press** 

Full text available: mpdf(1.80 MB)

Additional Information: full citation, abstract, references

In this paper, we describe a case study of the design and development of a group-

conferencing tool suite, built on top of an overlay network based event dissemination framework, which is extensible via quality of service template plug-ins. We explain, for each of the tools, how the framework built-in conveniences were explored to create simple but effective distributed solutions, backed by the appropriate quality of service templates, whose design we also discuss.

**Keywords:** active networks, case study, event dissemination, multicasting, overlay networks, quality of service (QoS)

Results 21 - 30 of 30

Result page: previous 1 2

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player